

CLAIMS

What is claimed is:

1. A multi-processor system comprising:
an owner predictor control that provides an ownership update message corresponding to a block of data to at least one of a plurality of owner predictors in response to a change in an ownership state of the block of data, the update message comprising an address tag associated with the block of data and an identification associated with an owner node of the block of data.
2. The system of claim 1, wherein the owner predictor control provides an ownership update message when the block of data at the owner node transitions to one of a modified or exclusive state.
3. The system of claim 1, further comprising a requesting node that provides a first request for the block of data to a home node, the requesting node being operative to provide a second request for the block of data to at least one predicted node in parallel with first request, the at least one predicted node being selected by an associated one of the plurality of owner predictors.
4. The system of claim 3, wherein the requesting node receives a coherent copy of the block of data from at least one of the home node and the at least one predicted node, the requesting node consuming a first coherent copy of the block of data received.
5. The system of claim 3, wherein a cached copy of the block of data exists at the owner node, the home node issuing a third request for the block of data to the owner node.
6. The system of claim 5, wherein the system employs a directory-based cache coherency protocol, the home node further comprising a directory that maintains directory state information associated with the block of data, the home node issuing the third request to the owner node based on the directory state information indicating that the owner node has an exclusive cached copy of the block of data.

7. The system of claim 5, wherein the owner node provides one of (i) a response to the home node and (ii) a response to the home node and to the requesting node, the owner node providing the response based on a state of the cached copy of the block of data at the owner node.
8. The system of claim 5, wherein the at least one predicted node comprises the owner node, the owner node having an exclusive cached copy of the block of data and providing a data response to the requesting node based on which of the second request and the third request arrives at the owner node first.
9. The system of claim 1, wherein a given one of the plurality of owner predictors, associated with a processor, comprises a first component that predicts an owner node of the block of data by observing the pattern of instructions within the processor and a second component that stores ownership update messages provided from the owner predictor control.
10. The network of claim 9, wherein the second component stores the provided update messages according to a first-in-first-out (FIFO) arrangement.
11. The network of claim 9, wherein the second component is operative to prioritize update messages according to a determination at the first component.
12. The network of claim 9, wherein the processor employs the given owner predictor to determine a predicted owner for a given block of data, the given owner predictor selecting between accessing the first component and the second component according to the frequency in which ownership update messages associated with the block of data have been received from the owner predictor control.

13. A multi-processor network comprising:
 - a first processor that includes a cache having a plurality of cache lines associated with respective blocks of data, one cache line transitioning to an ownership state based on a response to a request provided by the first processor;
 - a second processor that includes an associated owner predictor;
 - an owner predictor control that provides an update message to at least the owner predictor associated with the second processor to identify ownership for the one cache line consistent with the one cache line transitioning to the ownership state.
14. The network of claim 13, wherein the owner predictor control provides the update message to at least one other processor at a multi-processor node that is shared by the first processor.
15. The network of claim 13, wherein the owner predictor control broadcasts the update message to each of a plurality of processors comprising the multi-processor network.
16. The network of claim 13, wherein the owner predictor control monitors available bandwidth in the network and provides the update message based on the available bandwidth relative to a threshold value.
17. The network of claim 13, the network further comprising a home node having a directory that includes directory state information associated with the plurality of cache lines, the directory state information being updated to reflect the one cache line transitioning to the ownership state, and the owner predictor control providing an update message in response to the updating of the directory state information.
18. The network of claim 17, wherein the second processor provides a first request for data to the home node and a second request for the data at least one predicted node identified by the owner predictor..
19. The network of claim 18, wherein the at least one predicted node comprises the first processor based on the update message.

20. The network of claim 17, further comprising an unordered network interconnect that enables communication of requests, responses, and update messages among at least the first processor, the second processor and the home node.

21. A system comprising:

a requesting node that provides a first request for a block of data to a home node, the requesting node being operative to provide a second request for the block of data to at least one predicted node substantially in parallel with first request, the requesting node receiving at least one coherent copy of the block of data from at least one of the home node and the at least one predicted node;

an owner predictor associated with the requesting node programmed to identify the at least one predicted node; and

an update control that provides an ownership update message to the owner predictor in response to a detecting a change in an ownership state for the block of data, the update message comprising an address tag associated with the block of data and a processor identification associated with an owner node of the block of data.

22. The system of claim 21, wherein the at least one coherent copy of the block of data is returned to the requesting node as a response in a response channel, the response being provided by the at least one predicted node.

23. The system of claim 21, wherein the home node provides a third request for the data to an owner node if the owner node has an exclusive cached copy of the requested data.

24. The system of claim 23, wherein the first request is provided in a request channel, and the second and third requests are each provided in a forward channel.

25. The system of claim 24, wherein the system employs a directory-based cache coherency protocol in which requests in the forward channel do not block.

26. The system of claim 21, wherein the at least one predicted node comprises the owner node, the owner node providing a data response to the requesting node in response to which of the second request and the third request that arrives at the owner node first.

27. The system of claim 26, wherein the owner node provides a victim message to the home node and the data response to the requesting node in response to the third request arriving at the owner node prior to the second request, the home node providing a speculation acknowledgement to the requesting node in response to the victim message from the owner node.

28. The system of claim 26, wherein the owner node provides a victim message to the home node in response to the second request arriving at the owner node prior to the third request, the owner node also providing the data response to the requesting node in response to the second request from the home node.

29. The system of claim 21, wherein the at least one predicted node further comprises a target node having a cache that includes the data having one of an invalid state and a shared state, the at least one predicted node providing a miss response to the requesting node in response to the second request, and the owner node providing a data response to the requesting node in response to the third request.

30. A multi-processor system comprising:
means for identifying a predicted owner node associated with a block of data, the
means for predicting being associated with a requesting node;
means for selectively providing a first request for the block of data from the
requesting node to the predicted owner node; and
means for updating the means for predicting in response to a change in ownership of
the block of data, the means for updating being remote from the means for predicting.

31. The system of claim 30, further comprising:
means for providing a second request for the block of data from the requesting node
to a home node, the second request being provided substantially in parallel with the first
request; and
means for providing a coherent copy of the block of data to the requesting node in
response to at least one of the first request and the second request.

32. The system of claim 31, further comprising:

means for ascertaining whether the predicted owner node has an exclusive cached copy of the block of data; and

means for providing a third request for the block of data from the home node to an owner node when the predicted owner node has the exclusive cached copy of the block of data.

33. The system of claim 30, wherein the means for updating comprises means for determining a frequency with which the block of data has changed ownership over a period of time, the means for updating being operative to update the means for identifying for a the block of data based on the determined frequency relative to a threshold frequency.

34. A method comprising:

updating ownership state information for a block of data at a plurality of owner predictors associated with respective processors based at least in part on a change in the ownership state of the block of data; and

identifying at least one of the processors as a predicted owner node based on the updated ownership state information in an owner predictor associated with the at least one of the processors.

35. The method of claim 34, further comprising:

issuing a first request for the block of data from a requester to a home node;
concurrently issuing a second request for the block of data from the requester to the predicted owner node based on the updated ownership state information; and

receiving at least one coherent copy of the block of data at the requester from an owner processor, if the owner processor has an exclusive cached copy of the block of data, and from the home node, if no exclusive cached copy of the block of data exists when the home node receives the first request.

36. The method of claim 35, further comprising issuing a third request for the block data from the home node to the owner processor in response to determining that the owner processor has the exclusive cached copy of the block of data.

37. The method of claim 36, further comprising providing the at least one coherent copy of the block of data in response to the second request when owner processor receives the second request prior to the third request.

38. The method of claim 36, further comprising providing the coherent copy of the block of data in response to the third request when owner processor receives the third request prior to the second request.